Ask Mr. Protocol

by Michael O'Brien

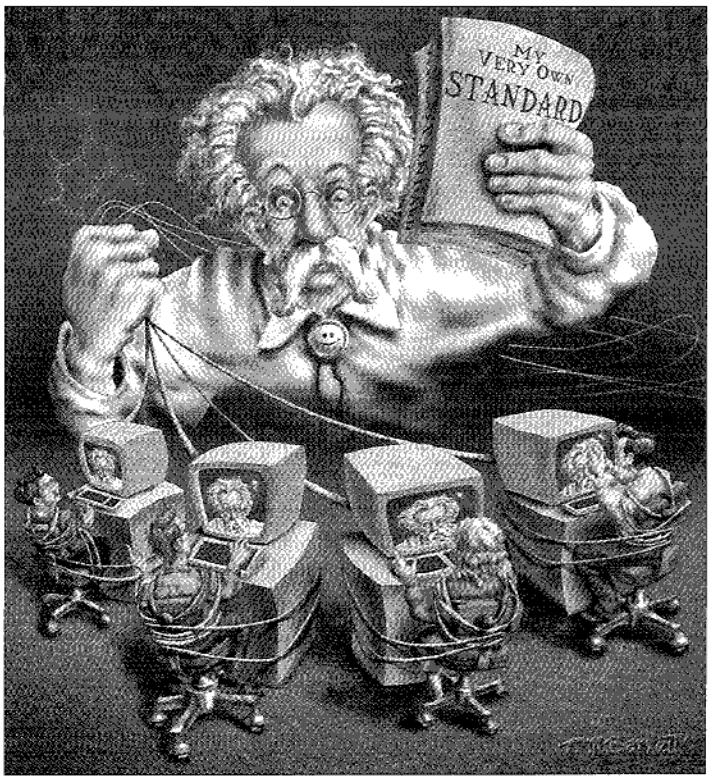

TOM BARRETT

"Security considerations are not discussed in this memo."
– Just about everybody

"Internet Drafts may disappear at any time."
– Rules of the game

"What do you mean we can't claim we're compatible? A standard's a standard, right?"
– Hungry vendor with acid stomach

Your Basic Standard Mr. Protocol

Q: *I have just developed the greatest protocol since the Angel of Death invented Passover. Thousands shall fall before the might of my checksums, and loud shall be the lamentations of those who fail to route my packets. The purity of my header bytes shall bring swift destruction and suffering unparalleled to my enemies. Death! Death and ingrown toenails to the infidels who fail to adopt it! And Victory! Victory unmitigated by sorrow, strife or packet loss to the faithful! ... Er... How do I get it adopted as a standard?*

A: No worries, mate. I'd say that you and the Internet Engineering Task Force are about ready for each other. You've got just the right attitude to engage in fruitful negotiations with these stalwart architects and guardians of the Internet.

Things in the IETF have never been what one would call placid. In a way, this is too bad: There have been more than a few people whose work would have made the Internet a better place, who have left the Internet standards arena because the interpersonal static in the IETF was too much for them. This has only gotten worse over the years, as the IETF population has increased, and as the IETF meetings have attracted a more diverse crowd.

As Mr. Protocol has noted before, the initial population of the Net included a much higher percentage of actual network implementers than today's. In addition, the network architects, designers and implementers were a more homogeneous population than they are today. Consequently, the path that protocols followed on the journey from inception to wide distribution was somewhat more informal than it is today.

The fundamental instrument for propagation of Internet standards has always been the Request for Comments. RFCs have a fixed format that makes it obvious that the network's bread and butter came from the U.S. government. Every single one has a required header reading "Security." This is there because the security folks didn't have too much trouble convincing the government and everyone else that network security has to be designed in from the beginning, not slapped on after the fact. That the network's initial design was done by academics is obvious from the fact that in almost every RFC, the sole contents of the "Security" section is the sentence: "Security considerations are not discussed in this memo."

Everyone who wanted to invent a protocol was free to do so. In fact, in large measure, this was what life was all about. Here was this nifty new network; let's invent some services for it! This was simple enough. First of all, you invented the protocol. Then you implemented it, and after you had a working implementation, you fixed it and fussed with it and futzed with it until you had the rough edges worn off, and then you wrote an RFC about it. You also talked about it on the relevant mailing lists and convinced someone else to write an implementation for some other kind of

machine. Then when the two implementations refused to work together correctly, you fussed with the implementation and refined the protocol and its description until they did, at which point you declared success and moved on. Except that people wouldn't let you, because like a Chinese person whose life you'd saved, The Protocol was now your responsibility and prospective implementers felt free to wake you up in the middle of the night to talk about pathological cases for the rest of your natural life.

Which is exactly the reward you'd prayed for, right?

The only tricky part here was the protocol number. In ARPANET days, you needed a socket number. For the Internet, you need a port number. Same thing, different name. It's a way of identifying just what user-level protocol this TCP packet is supposed to be. Every protocol on the Internet, from highest to lowest, has some sort of identifying number that makes it unique among protocols at that layer. User-level protocols like FTP, Telnet and the Web have port numbers assigned to them. Telnet is 21, FTP is 23 and the Web is 80. You can prove this with Telnet. Telnet to another machine's port 23 and you'll get an FTP prompt. Telnet to port 25 and you'll get an SMTP prompt (25 is the mail port). Telnet to port 80 and you'd better be ready to type some HTTP stuff. Below the user level, TCP and the Internet Control Message Protocol (ICMP), each have identifying numbers. IP itself is identified by a number, including a version number (currently 4, hence all the talk about "IPv4" and "IPv6," which is coming along to replace it).

The Internet Assigned Numbers Authority (IANA) assigns protocol numbers. In the early days, getting a new number was a simple matter of sending an email saying, roughly, "Gimme a number." You had to provide some information, such as the title of the protocol, your name and contact information, and a note from your mother saying that yes, this was a genuine protocol that people were actually going to use and not a master's thesis with a case of gigantism. And that was pretty much it. If IANA didn't hear from you for a while, you might eventually get a note saying, "Is this thing still real, or can we release the number now?" Aside from that, you pretty much owned the road as far as that particular protocol number went.

On the ARPANET, congestion was the only real problem the ordinary network user, or even the ordinary protocol implementer, could cause.

Now, it might be thought that if this were a network upon which people depended, allowing such free and easy experimentation might not be such a good idea. For a number of reasons, this cavalier attitude was not only allowed, but encouraged.

For one thing, that was the purpose of the net. The entire network was experimental, even though people depended on it. Eventually, the MILNET was split from the ARPANET, a year or two after the TCP/IP protocol suite replaced the NCP protocol, thereby allowing multiple networks with routing between them. It was this event that marked the first retreat from a regime of more or less untrammeled experimentation.

Second, at least on the ARPANET, it was hard to break the net. Practical knowledge of packet-routing algorithms was still a dicey proposition, but it wasn't a proposition that users ever encountered directly. The ARPANET was a two-level net. Only the Interface Message Processors (IMPs) communicated directly over the net, and they were homogeneous, centrally maintained and shared a common base of routing software. Users didn't have it in their power to alter or affect this level of the network in any way, aside from an unfortunate proof-of-concept escapade at Massachusetts Institute of Technology that partitioned the net for a day or two. Congestion was the only real problem the ordinary network user, or even the ordinary protocol implementer, could cause.

Not that they didn't try. Before the IP protocol was deployed, institutions could connect only one or two mainframes to the net, that being all that the IMPs could support. As the brand-new concept of local-area networking spread through the community, more than one institution seriously proposed various sorts of gimmickry to allow the host connected to the IMP to act as a gateway. NCP couldn't support this in any straightforward manner, there being no such thing as a network number in NCP. The schemes involved doing various sorts of violence to the host address. These ideas were rejected. As the users didn't actually have any access to the internals of the IMPs, that was the end of that idea.

Enter Internet Standards

As the Internet has grown, the "standards path" has become more elaborate. Initially, there wasn't one. Everything was an RFC, plain and simple. Some RFCs (the one specifying IP, for example) were more important than others (the one for obtaining time and temperature in the machine room). The fact that there were no standards was regarded by some as a strength, because it allowed people to try alternative solutions to established practice without sanctioning hordes of people crying out that "standards are being violated."

As more and more people's livelihoods came to depend on the Internet, and to make it easier to deal with international bodies–which believe that standards are fundamental–the Internet community eventually began to evolve its own standards track process. The RFC describing the process has undergone two revisions since it was first issued.

Some things remain the same. First of all, RFCs are not synonymous with standards documents. Not all RFCs are standards documents. Anybody can still write an RFC and submit it for publication. However, very few RFCs are now submitted out of the clear blue sky. A new type of document, the Internet Draft, has been created to allow work in progress to be published.

Internet Drafts are peculiar beasts. If left alone, they disappear automatically and tracelessly. Consider this a type of intellectual garbage collection. They

may be revised at any time, replaced at any time or pulled at any time. The "sense of the Senate" is that they may not be quoted in any paper, publication or report, and they may never be referred to in the sense of a standard to be complied with. The aim here is to prevent the Internet implementers from inventing some new protocol, and while the protocol is still in the fiddling and futzing stage, have some vendor suddenly and loudly announce that they support it. Of course this happens anyway, but in general, vendors don't refer to Internet Drafts directly, otherwise their own representatives in the standards process will be given a hard time (which they are pretty much guaranteed to pass along within their company).

The normal way in which a standard comes into existence these days is that a working group is formed within one of the "Areas" of the IETF. The IETF meets three times a year, and working groups commonly arise out of "Birds of a Feather" sessions held at these gatherings. If the BoF session decides that a working group is warranted, they work with the area director to form one.

Once formed, the working group creates a mailing list and begins work on defining the protocol specification. Working groups, although the fundamental bodies for the creation of new standards, are themselves nonstandard. Internet architects and implementers include in their number a fair proportion of colorful individuals, and people skills are not high on the requirements list for membership. If physical assaults are rare, it's only because these people see each other a maximum of three times a year.

How does any work get done? Mr. Protocol is glad you asked.

The mailing list, alternating with IETF meetings, generally gets the job done despite the occasional noise. Once in a while, a working group fails due to internal politics, and the ensuing stench is great. This is a remarkably rare event, all things considered. Generally, people either get tired of arguing, or get disgusted and leave, or (mirabile dictu) work out their differences until an Internet Draft is created specifying the protocol.

At this stage, the working group chair requests that the Internet Draft be entered onto the standards track as a Proposed Standard. The IETF proper doesn't have much of a structure aside from the area directors, so oversight is given to the Internet Engineering Steering Group, a much smaller body elected from the membership. The IESG oversees all advances along the standards track, and if it approves, it instructs the RFC editor to advance the draft to an actual RFC.

A Proposed Standard has to be pretty much complete. In particular, it's a big boost if it's actually been shown to work in the real world, without breaking the world in consequence. Once two independent implementations have been created and real-world hang-ups fixed, the IESG may order the specification to be advanced to Draft Standard. At this point, the main thing standing between the spec and life as a full Internet standard is a large amount of real-world experience, together with a show of popularity. At each stage of the specification's progress, a new RFC is issued. RFCs, once issued, never change. That way, RFC 2036, for example, is always going to be the same. It may not be current, but it will always have the same text.

How to tell what's current? By giving RFCs two names. There is a parallel collection of STD documents, which represent the current standards. The contents of these do change over time. The change is made by pointing the current STD label away from the old RFC and toward the new RFC. The resulting document has a new RFC number but the same old STD number. There is an STD document, periodically updated, which gives the current status of all documents on the standards track.

The third version of the RFC detailing the standards track makes interesting reading. A sizable chunk of it deals with a staged process of appeals, going up through the IESG and the Internet Activities Board all the way to the Trustees of the Internet Society. None of it talks about lawsuits, which is too bad, because if the bozos currently attempting to litigate their way to fame and fortune by suing the IANA were thereby to lose their standing in the standards process, life might be a great deal simpler...at least for the IANA (and probably, eventually, for the rest of us).

Standards don't last forever, and those whose usefulness has passed are retired to Historic status. As the standards track RFC points out, the proper word to use here would be Historical, but the use of the word "Historic" is now itself historical.

Room to Experiment

Finally, there are protocol specifications that are not on the standards track, will never be on the standards track and which basically should never come near the standards track. The bin for these things is labeled "Experimental," and if that word strikes terror into your heart, you've demonstrated good perception. This is where most of Mr. Protocol's own new protocols wind up. The protocols in this category are either intended for a narrow use, are driven by some other standards body or are just plain outright dangerous.

Consider a superset of TCP, designed for use with satellite communications. Now, ordinary TCP, although it doesn't say so in the spec, backs off and slows down whenever it detects packet loss because it assumes that in today's Internet, substantially all packet loss is due to congestion. Well, in the satellite world, most packet loss is just due to loss: When one bit in 10^6 or so goes flooey, you're going to get dropped packets and that's that. The correct response in this case is not to back off. If anything, it's to speed up for a while. Detection of the dropped packet requires that the network pipeline to the satellite be drained and refilled, so the faster the source can do that, the better. In fact, an experimental protocol for satellite transmission has been created that does exactly this but still behaves well in genuinely congested scenarios. For various reasons, this protocol has a wardrobe of different

faces it can put on, only one of which looks like TCP. Given all this, and given the fact that it's not designed to replace TCP globally, now or ever, you have a perfect candidate for "experimental" status.

Perhaps the most interesting feature of the IETF standards process, aside from the fact that it is continuously evolving as people discover new mud pits to fall into, is the degree to which it has managed to preserve the network's egalitarian beginnings. Although most of it is taken up with the humdrum stuff of working group creation and management of the group's output, it also documents a second path along the standards track for standards that are not the output of a working group.

These individual contributions have proven highly useful in the past, so a path remains for the individual to propose a new protocol, independent of the IETF process. The RFC editor may suggest that the individual share the work with a working group if that seems appropriate, but if no appropriate working group exists, and the individual cannot or does not want to pull together a working group, the document can move along the standards track like any other, as long as the IESG agrees to promote it.

It gives Mr. Protocol a warm and happy feeling to know that his own personal efforts still have a standards track outlet. ➸

Mike O'Brien *has been noodling around the UNIX world for far too long a time. He knows he started out with UNIX Research Version 5 (not System V, he hastens to point out), but forgets the year. He thinks it was around 1975 or so.*

He founded and ran the first nationwide UNIX Users Group Software Distribution Center. He worked at Rand during the glory days of the Rand editor and the MH mail system, helped build CSNET (first at Rand and later at BBN Labs Inc.) and is now working at an aerospace research corporation.

Mr. Protocol *refuses to divulge his qualifications and may, in fact, have* none whatsoever. His email address is amp@cpg.com.